

Enabling Structure

In the early 1970s, feminist political scientist Jo Freeman wrote a paper intended for her sisters in the women's liberation movement whose main message was neatly captured by its title: "The Tyranny of Structurelessness."¹ Freeman pointed out that feminist groups were not impeded by the excessive hierarchy and bureaucracy that often characterize enterprises created and managed mainly by men. But if women's groups had successfully avoided those dysfunctions, then why did they not have a better record of getting things done, of achieving the purposes to which their members were so deeply committed? The answer, Freeman suggested, was that having no structure can be every bit as debilitating as having too much. In her view, what was needed in feminist groups and organizations was not to more adamantly eschew hierarchy and bureaucracy but instead to invent, adapt, and learn to use well structures more consistent with feminist values.

Freeman's paper was controversial when it first was distributed, but it would not raise many eyebrows today. The experiences of contemporary team managers, whether in ideologically driven organizations such as worker

cooperatives or in traditional businesses and public agencies, affirm the wisdom and the generality of her message. These days, expert team managers focus much more on identifying structural features that can powerfully and efficiently facilitate teamwork than on tearing down existing structures in hopes that teams will thereby be "freed up" to accomplish their work unencumbered by organizational red tape.

It is true that traditionally designed organizations often are plagued by constraining structures that have been built up over the years to monitor and control the behavior of individual employees. Inappropriate or over-specified task structures, personnel policies, and control systems can indeed impede productivity when work is performed by teams. The question is what leaders should do in such circumstances. We saw in the previous chapter that it usually is futile to replace the exercise of legitimate managerial authority about team purposes with consensus decision making in hopes that a sharp and finely honed direction for a team eventually will emerge. It is just as futile to dismantle organizational structures in hopes of releasing a team's pent-up power. Leaders who do that often wind up providing teams with *less* structure than members actually need to accomplish their work. Tasks are defined only in vague, general terms. Lots of people may be involved in the work, but they may be the wrong people or there may be too many of them. Norms of conduct are left entirely up to the group on the assumption, as one manager told me, that "The team will work out the details." That is indeed what members will do—but, as Freeman noted, in the absence of an enabling team structure they may wind up wasting large portions of members' time and energy on interpersonal and political issues of little relevance to the team's main purposes.

Structure, like authority, is in itself neither good nor bad for teamwork. It all depends on the *kinds* of structures that are created. The best ones provide members with a solid platform on which to carry out their collective work but also leave lots of room for them to develop their own unique ways of operating. Rather than establish up front everything that may be needed for a team to perform well, wise leaders focus mainly on the handful of structural features that establish a good basic "frame" for the team's work and then give the team plenty of room to mold that frame to their particular circumstances.²

In this view, structuring a team has much in common with designing a house or office. Does the architect attempt to anticipate all the uses to

which the space may be put and then, in hopes of optimizing utilization of the structure, generate a design that seeks to direct and constrain the behaviors that will occur within it? Or does the designer recognize that all the ways occupants may want or need to use the space can never be known ahead of time, and therefore create a structure that is well tuned to the basic functions the space will serve but also incomplete, unfinished, and adaptable? The latter approach recognizes that organic forms are always in a state of development and are never fully finished, and is the architectural parallel of good team structure.³

When designing work teams, then, leaders should refrain from specifying too much—but they also should take explicit initiatives to put in place the basic structures that will foster team effectiveness and minimize the organizational obstacles to teamwork. The key to good team design is to differentiate wisely between critical and unnecessary structural features. The three structural features that our research has shown to be key in setting the stage for effective teamwork are the design of the work that the team performs, the core norms of conduct that guide and constrain team behavior, and the composition of the team. These three features are explored, in turn, in the pages that follow.

THE DESIGN OF WORK FOR TEAMS

Good work design for teams is a fairly straightforward extension of what has been learned about the properties of motivating individual tasks. Some years ago, Greg Oldham and I proposed a set of task attributes that foster what is known as *internal work motivation* for individual performers.⁴ A person who is internally motivated feels good when he or she performs well, and feels terrible when the work has gone poorly, thereby lessening the need for motivational props such as performance-contingent extrinsic rewards or close supervisory scrutiny. People have internal motivation when they view their work as meaningful *and* feel personally responsible for work outcomes *and* receive trustworthy knowledge of the results of their efforts.

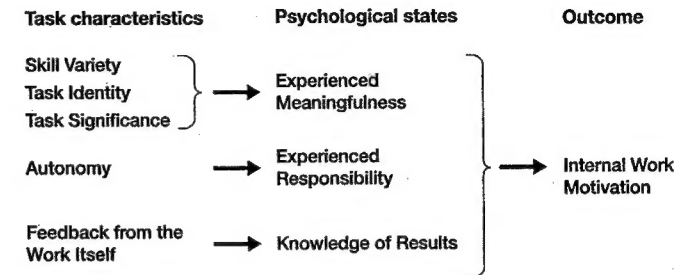
Even a task as inconsequential as writing a computer program to handle the routine management of one's personal finances can create the three psychological states and thereby elicit internal motivation. The task

is meaningful (at least to me: I'd love to rid myself of the tedium of bill paying, and besides, I find it challenging to write good computer programs). I have full responsibility for the work (I personally make all the decisions about program logic and coding). And knowledge of results is immediate and trustworthy (either the program runs correctly, or it does not). Take any of these three features away and internal motivation disappears. I cannot give myself a pat on the back if the work is trivial or entirely routine, or if I am not the one responsible for the work procedures (e.g., if I merely type in a program written by someone else), or if I submit the program but never see whether or not it runs properly.

What Oldham and I did was identify the measurable properties of jobs that give rise to the psychological states just described—experienced meaningfulness, felt responsibility, and knowledge of results—as is shown in figure 4-1. To illustrate, consider how the work of assembling a small appliance such as a kitchen toaster might be designed. On a traditional assembly line, a worker might do but a single and simple part of the overall task, such as attaching the power cord to the toaster chassis. A more motivating job, by contrast, would involve assembling the entire device, testing it, and perhaps even boxing it up for shipment to the customer. Making the toaster is *meaningful* to the worker for multiple reasons—she does the entire job from beginning to end, the work involves use of a variety of her skills, and it is inherently significant because the product will be valued by those who eventually use it. (“What did you do at work today, Mommy?” her child asks. There is a world of difference between saying “I attached lots of power cords” and saying “I made toasters for families to use in their kitchens.”) The worker feels personally *responsible* for the outcomes because she has considerable autonomy to make decisions about the work processes, rather than following to the letter a procedure that someone else engineered and that a supervisor enforces. And the worker has *knowledge of results* of the work, since she generates direct and trustworthy feedback by personally testing each toaster before it is shipped. Tasks that are designed in accord with these principles generally elicit far greater internal work motivation than do those that are simple, repetitive, and of little broader significance, that provide little or no latitude for decision making, and that rely more on supervisory assessments of how well the work has been done than on feedback built directly into the work itself.

FIGURE 4 - 1

Job Characteristics That Foster Internal Work Motivation



Source: WORK REDESIGN by Hackman/Oldham, (c) Reprinted by permission of Pearson Education, Inc., Upper Saddle River, NJ.

Might the same principles be applied as well to work that is performed by teams? Is there such a thing as *collective* internal motivation? The affirmative answer to those questions is obvious to anyone who has observed the differences in behavior exhibited by members of the winning and losing athletic teams immediately after a championship game, or by members of a project team who have just learned that their proposal has been accepted (or declined) by the client for whom it was prepared, or by a medical team that has just saved (or lost) a patient. The collective celebration, or shared gloom, that one sees on such occasions attests to the fact that internal motivation is just as real for teams as it is for individuals.⁵

What specific features of a team's work foster collective internal motivation? The answer is illustrated in the way the Butler Manufacturing Company structured the work of teams that manufactured large grain driers for its farmer-customers some years ago at the plant it operated in Story City, Iowa.⁶ Teams were given full responsibility for constructing entire three-story-high grain driers from start to finish. It was a complicated operation, involving thousands of different parts and five types of work: assembly, fabrication, machining, painting, and shipping. Individual team members moved among these different types of work, both to provide variety and to expand each team's repertoire of skills and its flexibility. Teams had considerable latitude in how they proceeded with their work, participated in planning and scheduling meetings to ensure that the flow of completed products would meet customer requirements, and had full access to plant

data about inventory, orders, and productivity. Teams also participated in the design and development of new products and tools and had the authority, in some areas and within specified financial limits, to purchase tools and materials on their own. Each team tested each drier it built, and then affixed to the device a sticker giving the team's name so that the farmers who purchased it would know who had made it and could answer questions about it. Team members even provided field service of the driers they had built. According to Larry Hayes, plant manager at the time, the field visits "teach them the impact on a farmer's business if a machine isn't working. They also learn more about the technical aspects: how the product is used in the field. And they become more in tune with our customers, our bread and butter."

The way Butler structured the task of making grain driers resulted in gains in all three of the criteria of team effectiveness discussed in chapter 1. First, the customers were well served by having driers that worked well—and by having direct access to the team that built their device if they encountered any difficulties. Second, the teams themselves became more competent performing units over time, as members learned together what assembly procedures worked best for them and became practiced in carrying them out. And third, individual team members experienced far greater personal learning from their teamwork than ever could have been obtained from performing any one part of the overall task. Given the high standing of the Butler teams on the three criteria of effectiveness, it comes as no surprise that the plant also scored high on traditional measures of quality and productivity and that profits were about 10 percent higher than comparable operations in other locations. In the words of the plant manager, "We really do produce a better product here."

The benefits of designing work for teams are substantial. But getting the team's work designed right often requires leaders to overcome powerful inertial forces in their organizations and, as will be seen in the following sections, always brings some measure of risk. Creating a teamwork-structure that fosters collective internal motivation is never as simple as just tossing the team a task and letting members run with it.

Bigger really is better . . . Designing work for teams makes it possible to create tasks that are large and significant. No single individual could ever construct a complete grain drier because it is far too big and

complex. But, as was demonstrated at Butler, a team could. No individual could possibly redesign the entire science curriculum for a secondary school because no one person knows enough. But a team could. No individual could handle all the requests for service that a computer reseller receives from a large corporate customer. But a team could. No individual could single-handedly clean the several floors of a medical clinic. But a team could.

For each of these pieces of work, those who structure the work have a choice. They could break the overall task up into pieces, assign each piece to an individual, and then devise means to coordinate the individuals' contributions and integrate them into a coherent whole. Alternatively, designers could assign the entire task to a team and give members the responsibility of coordinating their own efforts and switching subtasks as needed so that the whole piece of work gets done on time, efficiently, and in a way that satisfies the customer or client.⁷

The latter alternative almost always brings significant increases in the meaningfulness of the work. It is more meaningful to do the entire science curriculum rather than just the chemistry portion, or to be collectively responsible for all the service needs of an important customer rather than to be sent by a dispatcher from one user to another, or to share responsibility for the maintenance of an entire clinic rather than just vacuum the second-floor carpets day after day. The ability to create work that is challenging, complete, and significant—and therefore meaningful to those who carry it out—is one of the major advantages of designing work for teams.

. . . but social loafing is always a worry. As powerful as the motivational advantages of well-designed teamwork are, tasks designed for groups rather than individuals also bring significant risks. One of the most pervasive of them is the tendency of individuals to slack off when working in groups. Psychologist Ivan Steiner refers to this tendency as the "motivation decrement" and notes that it almost always is present to some extent in group work. Other psychologists call the same thing *social loafing*, and economists call it *free riding*.⁸

Just as a voter may ask, on a rainy election day, "What difference will my one vote make?" so may a member of a large work team ask how much difference it really will make to not come in to work today, or to leave unfinished work on the desk Friday rather than take it home to wrap up

over the weekend. Although the motivation decrement is more pernicious when the work itself is poorly designed, it is present even for team tasks designed in accord with good work design principles.

Experienced team leaders take steps to head off the motivation decrement not just by ensuring that the work is as well designed as possible but also by making the team as small in size as it reasonably can be. (Although bigger is better for team task design, we will see later in this chapter that smaller definitely is better for group composition.) And, once the work is under way, expert team coaches watch carefully both for signs of motivational slippage and for opportunities to help members develop and sustain high shared motivation (see chapter 6).

Autonomy gives teams room to excel . . . Collective internal motivation also is fostered by team tasks that provide members a large measure of autonomy to decide how they will use their human and material resources in carrying out the work. In terms of the authority matrix discussed in chapter 2, a work team should have at least the right and responsibility to monitor and manage its own work processes (that is, to be a self-managing team rather than a manager-led team). This is advantageous not only for motivational reasons but also because team members usually have a much better understanding of the demands and opportunities in their immediate work situation than do the managers or engineers who lay down and enforce standard work procedures. Standard procedures are almost always a little "off" from the ideal way to proceed in the particular circumstances of the moment, and they invariably result in members feeling less collectively responsible for how the work turns out than would be the case if the team were genuinely self-managing.

Teams whose tasks provide ample autonomy for managing work procedures also have the opportunity to experiment with alternative ways of working together and, through trial and error, to become more competent in their work than ever could be imagined by those who determine the "one best way" the work should be carried out. It is well known among those who do shift work that the midnight shift, despite its disruption of sleep cycles, can be the best time to work because there usually is no one around to enforce adherence to standard procedures. Workers on the midnight shift commonly devise their own, better ways of getting the work done, take full responsibility for the results they produce, and feel collective pride

in what they accomplish. We all should learn from the midnight shift, and give work teams in organizations the kind of autonomy about work processes that can engender a strong sense of collective responsibility among members and nurture their impulse to improvise, experiment, and learn from the team's successes and failures.

. . . but autonomous teams gone bad are very bad indeed. Work teams do not always use their autonomy in ways or toward ends that those who designed the team's work had in mind. Although a team with autonomy has the power to do wonderful things for its customers or clients, it also can do real damage. Recall, for example, how one of the flight attendant teams at the domestic airline (chapter 1) used its considerable autonomy to line the pockets of team members rather than to provide their customers with the best in-flight service of which members were capable.

Managers understandably seek to protect themselves and their organizations from the disasters that can be wrought by a work team that goes sour. Sometimes that protection is obtained at the cost of unexploited opportunities, such as Supervisor Szczarba's group of telephone service providers, described in chapter 2, who were a team in name only. Managers in that organization did not have to worry about the team going out of control because there was no team there. But neither was there any possibility that the service providers could develop into a superb work team. Another costly way to obtain protection from the damage that can be done by a rogue team is to strip the team of its autonomy by laying out in great detail all the work procedures members are supposed to follow. But by denying the team the latitude to adjust its performance strategies to deal with unanticipated problems or to invent ways to exploit emergent opportunities, this strategy trims the team's upside potential just as surely as it protects against downside risks.

Problems are especially likely to emerge when managers *tell* team members that they have the authority to manage their work and that they therefore are accountable for the results—but then specify work procedures in such detail that members have no way to actually exercise that authority. The response of the crew of the Skylab 3 space station to the restriction of its autonomy by planners and managers in Mission Control provides a vivid example of how a team can respond in such circumstances.

NASA rhetoric emphasized (indeed, sometimes came close to glorifying) the key role that the astronaut crews played in Skylab missions. Yet the actual behavior of the Skylab 3 crew was almost completely dictated by Mission Control in Houston, even to the extent of what crew members were to do during meal times. Here is how Neil Hutchinson, the lead flight director for the mission, put it:

Back at the first mission, we weren't good enough to schedule the guys tight, but by the time the second mission ended, we knew exactly how long everything took. . . . We knew how long it took to screw in each screw up there. . . . We prided ourselves here that, from the time the men got up to the time they went to bed, we had every minute programmed. . . . You know, *we* really controlled their destiny.⁹

The crew repeatedly communicated members' frustration and irritation with how Mission Control was managing the crew's time and work, but with no apparent effect. Finally, the crew took an action that definitively asserted their own control of the mission: They turned off the radio and refused to speak with Mission Control. It was the world's first strike in space.

Although the astronauts' action was extraordinary, it reflects a more general and pervasive phenomenon. When managers view the work of a team as especially important, they often find it nearly impossible to keep from specifying all procedures in detail so that things will be done correctly. That way, managers reason, they can lower the probability of disaster to nearly zero—but in doing so they also severely restrict the team's autonomy to manage its own work. As we saw for both aircraft flight deck crews in the previous chapter and the Skylab 3 crew in this one, that can undermine the very aspiration that the procedural specifications were put in place to achieve.

Feedback makes team learning possible . . . Learning requires knowledge of results. When a team task is structured so that trustworthy feedback about performance comes not just to individual members but also to the team as a whole, learning opportunities expand.¹⁰ If there are others with whom members can compare their own reactions to feedback, individuals can make better sense of it and more productively probe its implications for their own work activities. Moreover, if the team has

become a setting in which members feel psychologically safe to explore the reasons for team successes and failures, considerable *collective* learning can occur—learning that would be far less likely if the feedback were exclusively about the performance of individual members.¹¹

Collective learning from performance feedback cannot occur, of course, unless a work team is reasonably well bounded and stable over time (see chapter 2). There was no way that teams of flight attendants at the international airline discussed in chapter 1 could have learned from postflight feedback, since crew composition at that airline was constantly in flux. By the time a customer's comments reached the airline, the team that served that customer no longer existed. All the airline could do (and did do) was put a note in each member's personnel file, which may have been helpful subsequently in individual performance assessments but was irrelevant to the improvement of the flight attendant teams.

Collective learning is not automatic even when feedback comes to a reasonably stable team whose members share a sense of psychological safety. We will see in chapter 6 that one of the most important functions a team coach serves is to help team members work through the antilearning temptations that invariably accompany both successful and failing performances, and then harvest the lessons that always are present when a team receives feedback about its work performance.

When the work provides a stable team with regular, trustworthy feedback about how it is doing *and* when the team is well coached, then the team is almost certain to evolve into a self-correcting performing unit, one whose every experience may come to be viewed by members as an occasion for continuous improvement.¹² Unless a team has data about how it is doing, however, there is no way it can learn. And unless a team learns, there is no way it can improve.

. . . but a poor team design can foster an antilearning stance.¹³ Members of poorly designed teams sometimes ignore or deflect opportunities to learn and improve even when they do receive abundant feedback about how they are doing. I once worked with a social service team in a state welfare agency whose members spurned virtually all the opportunities they had to learn from the considerable feedback the team received—from clients as well as from agency-generated statistical reports and supervisory assessments of the team's productivity and efficiency. During the

several months I occasionally observed the team, members never used any feedback from any source to reflect on ways the team might improve its productivity, efficiency, or service quality. To the contrary, members constantly reassured one another that they were doing splendidly in a nearly impossible work situation. Managers who raised questions about their performance or suggested ways the group might improve were viewed as not understanding anything about the rigors of front-line work. Clients were characterized as people who just wanted "more" and who had no appreciation of the quality of the services they actually were receiving. And the statistical reports that the team received were seen as irrelevant if not actively misleading. In conversation with one another, and with me, members affirmed that they were a fine social service team, and they explicitly rejected or rationalized any data that would have implied otherwise.

This team was not composed of bad people. Most members were deeply committed to their work and their clients. They had accepted lower salaries and higher workloads than they would have had in other organizations because they believed so strongly in the importance of what they were doing. But the team was poorly designed—its direction was less clear than it could have been, and although members had a good deal of latitude in how they dealt with client families in the field, legal requirements and agency policies significantly constrained their autonomy to do what they felt was in their clients' best interests. Perhaps partly because of the flaws in its design, the team had fallen victim to a common group malady, the tendency to see one's own group as virtuous and other groups (in this case, both clients and management) as impediments to their own good efforts. Even when there is a great deal of feedback available, as there was in this case, teams that are poorly designed may develop an antilearning stance that closes members off from the very kinds of improvements that both they and those they serve would greatly value.

Balancing Benefits and Risks

There are many benefits, for organizations and for their clients or customers, of designing work to be performed by teams. As we have seen, team tasks can be structured so they challenge members and allow them to do an entire and significant piece of work from beginning to end, thereby boosting the experienced meaningfulness of the work. Teams can

be given substantial autonomy, within clearly specified limits, to decide about work methods and procedures, thereby enhancing members' sense of collective responsibility for work outcomes. And performance feedback can be provided to the team as a whole, which both gives members socially verified knowledge of the results of their work and the opportunity to learn from their collective successes and failures. When present, these features almost always create a state of collective internal motivation and, in many cases, can set a team onto a course of continuous improvement.

Risks also are introduced when teams perform work, however, which can negate or even reverse the upside potential of teamwork. This is as it must be. Although well-designed work is a key ingredient in setting the stage for team effectiveness, even the best possible work design cannot ensure good performance. Only when an enabling work structure is reinforced by engaging direction (chapter 3), a supportive organizational context (chapter 5), and expert coaching (chapter 6) are the benefits of good work design realized and the risks of designing work for teams minimized. How the work is structured is an important thing, but it is not the only thing.

CORE NORMS OF CONDUCT

Group norms specify what behaviors are acceptable—and unacceptable—in a group. Behavior that is viewed as appropriate by the team is reinforced, and behavior that is seen as unacceptable or inappropriate is sanctioned. Given that the approval and disapproval of one's teammates are consequential for almost everyone, a member's behavior can be shaped readily by almost any group of which he or she is a voluntary member.

One of the great things about norms is that they can be about anything members want, although in practice they tend to focus on behaviors that members view as especially important. So if team members decide that they want to have a norm that people do not interrupt one another, or that everyone will arrive punctually at each meeting, all they have to do is obtain agreement about that. Henceforth, members who interrupt others, or who are late for meetings, are likely to experience raised eyebrows or head shakes from their teammates. When a team norm is powerfully shaping behavior, one does not observe much deviant behavior (for the same reason one does not see people putting their fingers on a hot stove to make

sure it still burns like it used to). Moreover, the more members agree about what is approved and disapproved, and the stronger those sentiments, the greater the compliance with team norms.¹³

Team training courses in organizations and practical books about good team behavior invariably give a great deal of attention to the development and enforcement of group norms that are thought by the trainers or authors to facilitate effective teamwork. The norms mentioned earlier (about interrupting and punctuality) are commonly included in lists of "good" norms, as are norms about listening, information sharing, participation, respect, trust, and risk taking. I would much prefer to be in a group that has such norms than one that does not.

But does it follow that these norms are the ones that those who create work teams should explicitly build in as part of a team's basic structure? Let me propose that they are not, that norms such as these address behaviors that are secondary rather than primary enablers of team effectiveness. More fundamental are norms that are outward looking, that address the relationship between a team and its performance context. In their most general form, these norms are as follows:

1. Members should take an active, rather than a reactive, stance toward the environment in which the team operates, continuously scanning the environment and inventing or adjusting their performance strategies accordingly. (A team's strategy is simply the set of choices members make about how to carry out the work.¹⁴)
2. The behavioral boundaries within which the team operates should be demarcated, identifying the small handful of things that members must always do and those they must never do.

The first of these two norms lessens the risk that members will charge ahead with a work strategy that is inappropriate for the task being performed, or that they will fail to notice environmental opportunities or obstacles. The second norm lowers the chances that a team will get into trouble by inadvertently violating requirements or constraints established by its own organization. Teams whose behaviors are guided by these two norms are less likely to encounter unpleasant surprises in their work, and more likely to develop performance strategies well tuned to their task and

situation, than are teams whose norms exclusively address within-team interactions. Outward-looking norms facilitate competent team performance, to be sure. But they also provide a solid platform on which team members can develop, along the way, whatever additional norms they may find helpful in guiding and regulating internal team processes.

How the two basic norms that foster team effectiveness operate in practice is illustrated by the flight attendant team at the domestic airline that served the vacationers going to Florida on a sunny day and then, on a subsequent trip, staffed a weather-delayed plane full of business travelers trying to get to Boston (see chapter 1). Before each of these flights, the flight attendant team quickly scanned its work environment (the first of the two basic norms) and determined that different work strategies were required for the vacationers than for the business travelers. The strategies the team opted for could hardly have been more different for the two flights—a lighthearted and entertaining set of routines for the Florida trip, and a competent, low-profile, businesslike way of operating for Boston.

Team members also adhered to the "must do" and "must never do" prescriptions that had been laid down by their own management. On both flights, they performed each and every required safety procedure (the "must do" items). And on neither flight did they give away liquor (the "must never do" item) even though that might have further lightened the mood of the Florida vacationers and somewhat placated the delayed business travelers on the Boston flight. This flight attendant team had incorporated the two outward-looking team norms into its everyday work life, which enabled members to exploit the special opportunities the team encountered, sidestep potential difficulties (especially on the Boston flight), and do all that while staying well within the bounds of organizationally acceptable behavior. Moreover, on both flights the team adjusted and implemented its work strategy efficiently; a new flight attendant team—one that had not yet developed its operating norms—would have needed to invent its strategy from the ground up.

The flight attendants at the international carrier (discussed in chapter 1) also operated in accord with group norms. For that team, however, the norms were inward focused, aimed at minimizing disruption and confusion among members as they went about their work. Each crew member was expected to know his or her job and to follow exactly the choreographed moves that had been engineered to give the appearance of a seamless team

performance. These norms were shared across the entire flight attendant workforce, which was necessary because teams did not stay together long enough to develop their own unique ways of operating.

Although a flight attendant could join any team without worrying about his or her teammates' expectations (they were basically the same from team to team and from trip to trip), the international carrier's inward-looking norms did not encourage teams to scan their work environments and tailor their performance routines to each flight's unique characteristics. Only when one team's aircraft was diverted to an unfamiliar airport and members had to find meals and accommodations for a planeload of passengers did the team break stride. And it was only then, when standard team norms had become irrelevant, that the team became outwardly focused and my colleague on the flight was able to see the level of motivation and ingenuity of which members actually were capable.

Countering Ordinary Human Tendencies

It would be convenient if the two core norms that foster team effectiveness appeared automatically, but unfortunately they do not. Instead, they almost always must be deliberately created as a feature of a work team's structure. The reason they must be explicitly created is that they operate in opposition to a pair of very ordinary human tendencies that can impede competent teamwork.

Reacting to Whatever Comes Our Way. The first ordinary tendency is our disposition to react to whatever captures our attention and demands a response, rather than to actively scan our environment for less obvious problems and opportunities that may call for nonstandard actions. It is a stimulus-response kind of thing: A stimulus comes our way, and we respond. A student sends me an e-mail asking for an appointment (the stimulus), and I reply by suggesting a specific time we could meet (the response). Another student writes to ask me to read the draft of a research paper, and I immediately agree to provide comments. Over time, my responses become habitual: I reply without explicitly considering the request or contemplating alternative responses I might make. It is this pattern that time management books seek to break by asking readers to be clearer about their priorities, more aware of features of their environments

that impede progress toward their main goals, and more deliberate about developing a behavioral strategy that really does put first things first.

The same things happen for task-performing groups. Members develop collective routines for handling common events in group life and then mindlessly execute those routines when one of those events appears. Ask members of a team with strong norms why they behave as they do and they probably will respond with something singularly uninformative such as "I don't know, that's just how things are done around here." Translated, that means the team norms are working and have become both routinized and invisible.¹⁵

Let me indulge myself with an example from the academy. I've participated in numerous faculty meetings in my career, and few have surprised me. Certain standard items will be brought up for consideration, we will discuss some of them at excessive length, and then we will dispense with the rest in rapid-fire order. Rarely, if ever, do we stop and reflect upon any opportunities or constraints in our environment that do not come packaged as an agenda item. A norm that supported active scanning of our university and intellectual context, and explicit planning of our strategy for further developing our department, could be very helpful to us. But we are faculty members. We do not need such artificial props. So we just keep on keeping on, letting opportunities lie unexploited and problems we could have known about surprise us.

Seeking Harmony. A second human tendency that makes it necessary to explicitly build the two core norms into the structure of a work team is our understandable impulse to have harmonious interactions with others, to be approved rather than rejected by our teammates, and generally to keep anxieties as low as possible.¹⁶ This tendency sometimes leads us to thoughtlessly do things that perhaps we should not do, and to go further than we ought to go in pleasing our fellow team members or our clients. By building into the structure of a work team an explicit norm that sets the outer-limit boundaries of what behaviors are acceptable, team members can have both a clearer sense of the location of the line they should not cross and a collectively enforced agreement that that line will be respected.

Building the two core norms into the structure of a work team is especially important when that team is likely to be subjected to either (1)

a barrage of demands that must be responded to quickly (in which case members may feel that there is no time to do *anything* other than respond as rapidly as they can) or (2) strong political or social pressures to cross the team's outer limits of acceptable behavior to please a customer, client, or supervisor (in which case team members may be strongly tempted to abandon their own standards in order to gain the approval of valued others).

Both of these conditions were in place for the fiscal analysis teams at OMB discussed in previous chapters. These teams had to deal with an almost unending stream of incoming questions and requests. It would have been the easiest thing in the world for team members to fall into a pattern of reactivity, dealing as competently and promptly as they could with all that came their way and giving themselves a pat on their collective back when, at day's end, their collective inbox was empty. For manager David Mathiasen, that was not good enough. So he reinforced both of the core norms whenever he had a chance. He made it clear to fiscal analysis teams that although responsiveness to client requests was important, it was not the only thing—or even the most important thing—they should be doing. The most important thing, he emphasized, was making sure that the teams' clients, most of whom were political appointees, were as fully informed as possible about the budgetary implications of the president's policies and programs. And that meant that the fiscal analysis teams needed to have their antennae out at all times so that members could be aware of what was developing in the background and be on top of emerging issues well before they spawned urgent client requests.

Second, as we saw in chapter 3, David set direction for the fiscal analysis teams in a way that firmly established the bounds of acceptable behavior. Analyses were to be conducted and results reported at the highest level of professional competence—the “must always do” item. Moreover, teams were never to bow to social or political pressure from their clients or from anyone else—the “must never do” item. The clients, David noted, also had to deal with their own stream of urgent demands and, as political appointees, were appropriately partisan. For both of those reasons, they could not be expected to insist that analyses be as comprehensive and objective as possible. Responsibility for that, David said, lay squarely with the civil servants who comprised the analytic teams. If a client wanted to truncate a chart so that the long-term budgetary impact

of some policy would not be seen, or stretch the y-axis so that an effect would seem especially large, that was the client's business. These were not activities in which fiscal analysis teams should participate.

How Norms Form

The core norms David established worked well for the fiscal analysis teams at OMB. Teams rarely were surprised by political or economic developments and they usually were ready with a strategy for dealing with emerging issues even before clients asked about them. Would these norms have developed naturally, without having been explicitly created by the teams' manager? To answer that question, we must take a brief detour to understand the three different ways that group norms can come into existence.

The first way group norms can be established is for them to be “imported” to the group by individual members.¹⁷ Each member brings to the group, based on previous group experiences, a set of reasonably well codified expectations about the kinds of behaviors that are acceptable in groups of this kind. If members' experiences have been similar, behavior will unfold in an orderly fashion, guided by invisible norms that never are explicitly discussed. This happens all the time. We all know what one is supposed to do and not do in a wide variety of group settings, ranging from a training seminar to Thanksgiving dinner. Behavior generally is orderly in such settings without group norms ever being discussed.

A second way that norms can become established is for them to evolve gradually, as members try out different behaviors. Invariably, members discover that some of those behaviors work well and are positively valued in the group, whereas others do not have the intended effect and may even get the actor into a little trouble with his or her peers. Over time—and generally it is a relatively short time—norms come into being for those behaviors that members view as important. Some time ago, I was asked to join a diverse group of faculty and administrators to conduct a performance review of a university staff member whose contract was up for renewal. None of us had done such a thing before. We not only had to figure out a strategy for carrying out the review, but we also had to establish a set of norms about how we would work together. Although we never

explicitly discussed what our norms should be, by our second or third meeting a set of shared expectations about appropriate behavior had evolved and was invisibly guiding our individual and collective behavior.

The third way to establish norms is to deliberately create them as part of the group structure, as was the case for the core norms of the fiscal analysis team at OMB. This third mechanism brings us back to our original question. Might the core norms that foster team effectiveness come naturally into being through importation by individual members or through gradual evolution in the normal life of the group? Or is the only viable way to establish these two outward-focused norms to explicitly create them? I know of no direct evidence that answers this question, but a considerable body of knowledge about behavior in social settings indirectly bears on it. That evidence suggests that it is quite unlikely that members will establish norms that support active environmental scanning and strategy planning, or that they will explicitly set and enforce specific "must always do" and "must never do" constraints on the team's task behaviors. Instead, the norms that members import or evolve are much more likely to focus on keeping interpersonal relations within the team and with clients smooth and conflict free, on keeping members' anxieties low, and on making sure that all inputs received are converted into outputs in a timely fashion with a minimum of fuss.¹⁸

The norms that are most likely to develop naturally, then, are likely to be good for maintaining harmonious interpersonal interactions within the group but less likely to be what is needed to foster team effectiveness and the long-term well-being of those the team serves. The two core norms are unnatural, and the behaviors they support often raise rather than lower anxieties within a work team. And that, in the final analysis, is why they usually must be explicitly and deliberately created as part of the team's structure.

How Secondary Norms Help

If the two core norms are in place, there remains a great deal of room for a work team to develop whatever secondary norms members may find helpful in guiding and regularizing their interactions with one another. I am entirely agnostic about what those norms should be, and it is none of my

business in any case. Secondary norms properly cover whatever behaviors members wish to regulate, and that is certain to vary from group to group.

Certainly many groups will find it helpful to have norms about such matters as punctuality, participation, communication, and conflict management. One insurance company, for example, printed up and distributed throughout the organization an attractive card that listed a handful of norms (e.g., "accept and share responsibility and participation," "never withhold information," and "never tolerate turf issues") that senior leaders believed would foster good interpersonal and group behavior throughout the company. That is fine, but it also would have been fine if their norms had focused on standards of politeness, or even on appropriate dress, if those were matters that members cared about. Secondary norms can and should be about whatever behaviors members view as important enough to regulate.

Far more significant for the health of a group than the content or coverage of secondary norms are the benefits that accrue to a group as these norms are created and enforced. In most groups there will be quite a bit of variation in the behavior of members early in the team's life (the exception being when all members of a group have such similar experiences that they import virtually identical expectations about appropriate behavior). As members find that they wish to rein in some of the more excessive, disruptive, or unhelpful behaviors that occur, group norms begin to form. As these behaviors come under normative control, members learn what their limits really are—what they want to regulate and what they are willing to let vary freely, which is less likely to happen when everyone always behaves well within the bounds of acceptability.¹⁹ Such knowledge is good for a team to have because it increases the likelihood that members will respond competently if the team subsequently has to deal with a problem for which standard ways of operating are inadequate or inappropriate. If a team has already established and tested its limits, it can more competently make decisions about whether and when to violate them.

Moreover, many groups include someone who regularly deviates from team norms. Because what that person says or does is so frequently "off" from mainline group behavior, he or she may gradually become identified as the group deviant. "There Vernon goes again," his teammates say, "always wanting to slow down the work until management shows it is taking our

complaints seriously." Vernon actually is serving the group well. In addition to clarifying the limits of what the team views as acceptable behavior, he is providing a corrective to majority processes that drive efficient task execution but rarely prompt innovation or change. By constantly having to deal with Vernon, his team is much more likely to come up with an original way of construing its situation, or new ways of responding to old problems, than would be the case if he also behaved well within the team's latitude of acceptable behavior.²⁰ Deviations from group norms, and the people who so deviate, serve their teams in important ways—even though their teammates may often wish that they would shut up, get in line, or go away.²¹

Getting Priorities Right

Norms are a critical aspect of team structure because they are a powerful and efficient means of coordinating and regulating member behavior. But not all norms are equally helpful in promoting team effectiveness. A leader's first priority should be to help a team establish the core norms that foster competent task performance—those that guide how the team deals with the opportunities and constraints in its performance context and that establish the outer limits of acceptable task behavior. Having done that, the leader can then stand back and let whatever secondary norms members may find helpful emerge as the team goes about its work.

I am aware that what I am espousing here runs counter to what most of us would do when seemingly "bad" behavior is exhibited in a team. Almost everyone would choose to focus first on getting that behavior corrected—tamping down whatever seems most disruptive and bringing those who are misbehaving back into line. And it is true that genuinely out-of-control behaviors within a team can make it impossible to get anything at all done, as would be the case if no one listened to anyone else or if most behaviors were aimed at putting other group members down rather than getting on with the team's real work. It is asking too much of a team to suggest that members develop outward-looking norms that may eventually help the team perform well when there is chaos inside the group.

In the great majority of cases, however, things are not that bad. In the normal course of teamwork, to give highest priority to dealing with unwelcome behaviors by establishing and enforcing norms that inhibit them is merely to treat the behavioral manifestations of what most likely is a more

basic structural malady. And it diverts members' attention and energy away from getting in place those core behavioral norms that actually will turn out, over the longer term, to facilitate achievement of the team's purposes.

The more powerful and enduring strategy for fostering competent teamwork is to focus first on ensuring that the basic design of the team is solid, and then help team members take full advantage of the positive potential of that design. The two core norms of conduct discussed in this section are a key feature of a good basic design. When these norms are in place, the probability that "bad" behaviors will occur is significantly reduced, and, moreover, team members are more likely to be able to deal competently with such behaviors when they do appear. Those who design and launch teams, therefore, should be insistent about the core norms, give the team great latitude to establish whatever secondary norms members may prefer, and then help the team learn how best to capture and use well the benefits of the "deviant" behavior that remains.

COMPOSITION OF THE TEAM

The three biggest mistakes people make when they compose a team are as follows:

1. They assume "the more the better" and therefore put too many people on the team.
2. They assume that people who are similar to one another will get along better, and therefore compose a team that is too homogeneous.
3. They assume that everyone knows how to work in a group, and therefore pay too little attention to the interpersonal skills of prospective members.

The mistake that they generally do *not* make is to overlook members' task-related knowledge and skill. We all know that a team is likely to get into trouble if members are not expert in the technical aspects of their work, so those who create teams generally take care to ensure that the team has plenty of task-relevant talent. Because that is the case, I focus here mostly on the three compositional errors just listed, and on what can

be done to avoid them, rather than on the necessity to make sure that team members actually know how to do that which they are supposed to do.²²

Size

It takes four people to play a string quartet, two crew members to fly a Boeing 737 aircraft, and twelve persons to form a full-sized jury. Not a person more nor a person less will do, so those who compose such groups can focus on matters other than the size of the performing unit. More commonly, however, managers who create work teams in organizations have considerable discretion about team size. Although managers sometimes form teams that are too small to accomplish their work well, the far more common and dangerous mistake is overstaffing them.²³

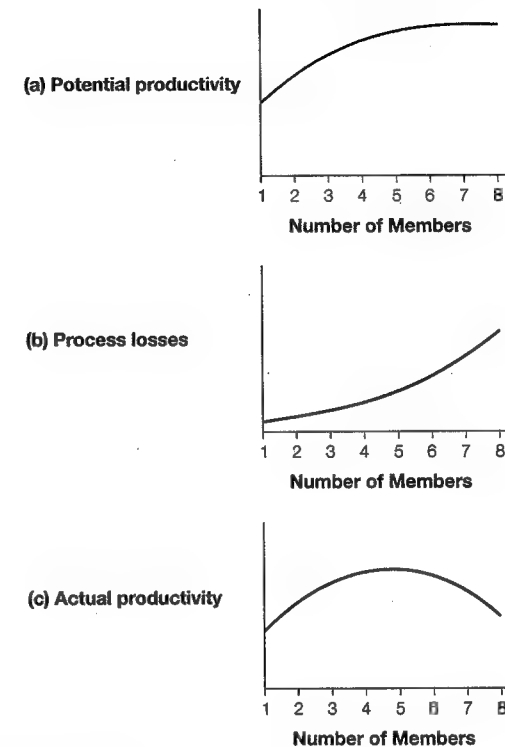
Frederick Brooks was the manager of the systems programming effort at IBM that in the 1960s created OS/360, then the largest such effort ever undertaken. As almost always happens with large-scale projects that involve a great deal of coordination and uncertainty, the project tended to fall behind schedule. The temptation in such cases is to compute how far behind the project is and then add staff to make up time. So if the project is a dozen person-months behind, perhaps a dozen people could be assigned to it for one month to get back on schedule. That has no better chance of working in software development, Brooks says, than would a scheme to produce a baby quickly by assigning nine women to be pregnant for one month each. In fact, adding people has the opposite effect, which led to the formulation of Brooks's Law: "Adding manpower to a late software project makes it later."²⁴

Psychologist Ivan Steiner reached the same conclusion in his analysis of the effect of group size on group productivity.²⁵ As can be seen in the three graphs in figure 4-2, the potential productivity of a group (that is, what the group theoretically could produce if member resources were used optimally) increases as size increases—but at a decreasing rate (figure 4-2a). Each new person adds something, but not as much as the previous person added. So increasing group size from two to three makes a much bigger difference in potential productivity than adding a thirteenth person to a twelve-person group.

Groups never perform at their level of potential productivity, however, because of what Steiner calls "process losses." These include the

FIGURE 4 - 2

The Relationship between Group Size and Productivity



Source: From Steiner (1972), p. 96.

motivational decrement discussed earlier, coordination problems, and the myriad of other inefficiencies that develop when people work together in teams. As is seen in figure 4-2b, process losses also increase with size—but they grow at an accelerating rate.²⁶ The actual productivity of a group, then, is its potential productivity minus the process losses. As is shown in figure 4-2c, actual productivity increases for a while as size increases, then levels off, and then actually begins to *decrease* for very large groups. When group size becomes very large, the problems generated far outweigh the incremental resources brought by the additional members.

Is there an optimum team size, one that maximizes the human resources the team has for use in its work without running the risk of the team getting bogged down trying to coordinate and maintain the engagement of a large number of people? A study conducted by Neil Vidmar and myself is at least suggestive of an answer.²⁷ We composed groups that ranged in size from two to seven members to assess the impact of size on group process and performance for various kinds of intellectual tasks. After the groups had finished their work, we asked participants independently to indicate the extent of their agreement with the following two questions. (You might wish to respond yourself for a group of which you are a member, using a scale from 1 ["strongly disagree"] through 3 ["neither agree nor disagree"] to 5 ["strongly agree"].)


1. This group was too small for best results on the task it was trying to do.
2. This group was too large for best results on the task it was trying to do.

We charted the average answers to these two questions on the same graph; the results are shown in figure 4-3. Not surprisingly, few people in the dyad thought it was too large, and few people in the seven-person group thought it was too small. What is noteworthy is where the two lines cross. We dropped a perpendicular line from that point to the horizontal axis on which group size was indicated (the dotted line in the figure), and voilà! we discovered that the optimum group size was 4.6 members.

That conclusion, of course, was just an exercise done on data from a not-very-important study, but it does remind us that most of the time smaller really is better. Indeed, a team may function better when it has slightly *fewer* members than the task actually requires.²⁸ A few pages ago, I noted that flying the Boeing 737 aircraft is a two-person task. In fact, Boeing engineers designed the cockpit so it could be flown by either a two- or a three-person crew. United Airlines was the lead customer for the aircraft and wanted the flight deck set up for a two-person crew, thereby generating enormous savings in labor costs. The pilots union, on the other hand, thought the plane should have a three-person crew since it would be flying many short hops in busy airspace. It surely would be safer to have a third pair of hands to help with the work and a third pair of eyes alert for potentially conflicting traffic.

FIGURE 4 - 3

Member Satisfaction with Group Size for Smaller and Larger Groups



Source: From Hackman and Vidmar (1970), p. 48.

When it became clear that the disagreement could not be resolved through discussion, the company and union jointly sponsored a research study to compare the behavior and performance of two- and three-person crews in actual flight operations. You no doubt have guessed the finding of the study: The three-member crews showed no across-the-board advantage relative to the two-person crews. Members of three-person crews did leave the cockpit more frequently to visit the cabin, which may have helped strengthen the work relationship between pilots and flight attendants. But they caught no more potentially conflicting traffic called to their attention by Air Traffic Control than did the two-person crews.²⁹

So what is the best group size? It depends on the size of the task, of course, but I do have a rule of thumb that I relentlessly enforce for student project groups in my Harvard courses: A team cannot have more than six members. Even a six-person team has fifteen pairs among members, but a seven-person team has twenty-one, and the difference in how well groups of the two sizes operate is noticeable.³⁰

If the evidence is so strong that small team size is better, why do we

see so many large teams struggling along in organizations? Certainly the faulty assumption that "more is better" for team effectiveness is part of the reason. But the main driver may have less to do with team performance than with emotional issues, such as using large numbers of people to share responsibility and spread accountability, and with political considerations, such as ensuring that all relevant stakeholders are represented in the group so they will accept its product. For these reasons, individuals from various constituencies may be appointed to a team one by one, or even two by two, creating a large team, a safe team, a politically correct team—but a team that can find itself incapable of generating an outcome that meets even minimum standards of acceptability, let alone one that shows signs of originality.

I once asked the executive director of a large art museum what in the world his forty-member board of directors could possibly accomplish. "Not much of anything other than to make financial contributions," he responded with a smile. "And that's just the way I want it." Sometimes those who create teams that are too large to perform well know *exactly* what they are doing. But what if one really does want one's board of directors (or top-management team, or some another team whose work requires many members) to be an effective performing unit?

There are a number of options.³¹ My university has a thirty-person Board of Overseers from whom are sought (and genuinely wanted) ideas, perspectives, insights—and, of course, contributions. But the Overseers are not the governors of the university. The decision-making group is the Harvard Corporation, which consists of five outside members plus the president and treasurer of the university, just the right size to make the consequential choices and decisions that are the responsibility of any organization's directors.³²

Another example. A startup organization I once studied became sufficiently large that the dozen founding officers of the firm no longer were able to coordinate their activities and make decisions in the informal manner that had worked well in the organization's early days. The time had come, the CEO concluded, for some kind of senior management structure. He considered devising such a structure himself, but that would have been inconsistent with the collegial and democratic spirit of the firm, which he and his colleagues greatly valued. He also considered the opposite strategy—asking the dozen officers to go off by themselves and

come up with a structure that they all could accept. But he wisely realized that a group of that size, all of whose members had enormous personal stakes in how things turned out, was quite unlikely to come up with a structure that would serve the organization well. What the CEO finally did was form a reorganization task force composed of four diverse members, all highly respected by their peers, whose task was to develop a proposal for the new management structure that he would review, possibly modify, and approve. But he did not stop there. He also created a "must always do" norm for the group. Each member was assigned responsibility for staying closely in touch with each of three other officers who were not on the task force. Before each task force meeting, each member was expected to speak to each of those three persons. The first agenda item at every meeting was to be a report of the views of those nonmembers, and the last agenda item was to be an explicit review of what should be communicated to, or asked of, them. Task force members took their responsibilities as linking pins seriously, and although there were many rough spots along the way (including slippage in adhering to the norm about pre- and postmeeting communication), the team eventually came up with a reorganization plan that was accepted both by the CEO and the other officers.

One aspect of Microsoft's strategy for protecting that company from the liabilities of large size has been to create organizational devices that make it possible for core programming teams to be quite small—commonly one program manager and three to eight developers. The company takes a modular approach to large development projects, providing each team with a clear and concise statement of the vision for its part of the work, specifying a clear deadline by which the work is to be done, and then giving each team substantial autonomy for doing what needs to be done to complete its part of the project. Even the development centers within which the teams operate, and whose managers watch over the links among the separate modules, are relatively small—usually no more than 300 to 400 people for work that might be performed in a 1,000-person unit in a traditionally structured organization. According to MIT management professor Michael Cusumano, from whose research this account was taken, Microsoft has found a way to make large teams work like small teams.³³

One final example. From its founding in 1972, members of the twenty-six-person Orpheus Chamber Orchestra have been committed to rehearsing and performing the repertoire for small orchestras in chamber

music style—that is, without a conductor and with the greatest possible participation by all members in matters of artistic interpretation (see chapter 6 for more on the orchestra). But a twenty-six-person team is far too large to operate as collegially as a string quartet does. With everyone chiming in with thoughts and ideas, rehearsal could become a cacophony. So orchestra members came up with the idea of the “core,” a small group consisting of the principal players for the piece being rehearsed. The core meets prior to the first full-orchestra rehearsal to work out the basic frame for the piece being prepared. Then, when the rest of the orchestra joins in, these individuals have special responsibility for helping other members of their sections understand and implement the ideas the core has roughed out. Any musician still can offer up new musical ideas for consideration by the ensemble, of course, but the starting point is the interpretive direction the core has set.³⁴

I have floated the idea of the core with players, conductors, and managers of a number of full-sized symphony orchestras. If it worked so well for Orpheus, could not the same idea be adapted for a 100-person orchestra? Could not principal players meet separately with the conductor before the first full rehearsal to work through the piece being prepared? It was a modest proposal, I thought, something at least worth thinking about if not experimenting with. But absolutely no one nibbled. It would violate the labor contract, I was told. Conductors would never stand for it. Players would resist. So large orchestras continue as they always have, playing great music to be sure, but doing so in a way that leaves enormous amounts of musical talent unused on the rehearsal stage and sufficing with less engagement and commitment from musicians than they could have.

With size, as with all other aspects of team structure, there always is a choice. But it takes the courage of informed conviction, plus a good measure of willingness to innovate and experiment, to find ways to exercise that choice that can simultaneously harvest the diverse contributions of team members *and* foster efficient collective action.

Mix

A well-composed team strikes a balance between having members who are too similar to one another on the one hand and too different on the other. Members of an excessively homogeneous group may get along well

together but lack the full complement of resources needed to perform well. An excessively heterogeneous group may have a rich diversity of talent and perspective but be unable to use it well because members are too different in how they think and behave. In a balanced group, members have a variety of talents and perspectives, yet are similar enough that they are able to communicate and coordinate with one another competently.³⁵

A balanced group is easier to describe than it is to create and maintain. Although the dysfunctions of too much diversity are real enough, the more common and pernicious problem in work organizations is excessive homogeneity. Multiple forces, acting in concert, foster similarity among members. For one thing, the people who are attracted to a given group are likely to have numerous attributes in common. “Which project would you most like to work on?” we ask. “And who would you like to work with on that?” The people who gravitate toward a particular kind of work, or toward a particular set of coworkers, are likely to be far more similar to one another than people who prefer different projects or teammates. Given a choice, I probably would be more comfortable working on an intellectual project with other white male pipe-smoking academics of a certain age whose first language is English and whose origins are Midwestern than with teammates who differ from me on all of those attributes. Such self-selection biases often are affirmed and extended by managers who understandably prefer to compose groups whose members, in their view, are likely to get along well together. And when team members have the authority to decide who will be invited to fill an open slot on the team, similar-to-us biases can dominate the selection process—at least for teams that are not yet mature enough to appreciate the value of diversity in accomplishing collective work.³⁶

Forces toward homogeneity continue to operate even after a group is formed and under way. Teams gradually but inevitably develop a shared view of reality (which means that members’ perceptions of the world in which they operate converge over time) and, as discussed earlier, they import or establish norms to regulate member behavior (which means that members’ behaviors usually become more similar over time). Once that happens, members’ beliefs and attitudes—which have been significantly shaped by the team’s shared perceptions and behaviors—gradually converge as well.³⁷ Finally, at some point in the life of most teams, one or more members decide to leave or are asked to. These individuals are not

drawn randomly from the membership; those most likely to depart are the ones who in one way or another are viewed as, or view themselves as, different from the rest of the team.³⁸

All of these forces pull in the same direction, and the result commonly is a team of rather similar people who find it easy to work together smoothly and harmoniously.³⁹ (We are, after all, the same kind of folks. Heck, maybe we could even have our next meeting on the links, or over a sherry down at the club.) Moreover, homogeneous teams are virtually certain to find it easier to develop the "shared mental models" that are all the rage these days in explaining what it takes for a team to perform effectively.

This is all to the good except for two troublesome facts. First, there is little evidence that homogeneous teams, for all the good relations their members may have with one another, perform better than heterogeneous teams.⁴⁰ Second, the experience of working in a homogeneous team is less likely to promote member learning (or, for that matter, learning by the team as a whole) than are experiences in a more diverse team. How much do I actually have to learn from that group of pipe-smoking white professors in whose company I am so comfortable? Not much.

In fact, the kind of task-focused conflict that is common when a diverse set of members find that they are not of one mind about how to proceed can improve decisions and even increase the chances that they will come up with something creative.⁴¹ Diverse groups do often experience rough sledding early in their lives as members struggle to figure out how to work together (and *interpersonal* conflict definitely is not a boon to performance).⁴² But if members make it through their early difficulties, they are likely to come up with products that are significantly more creative than those generated by more homogeneous teams whose interactions are smoother from start to finish.⁴³

The key to having a good mix of members is to balance carefully between too much similarity and too many differences. At Butler Manufacturing, the process began well before grain drier assembly teams were formed. Prospective team members initially were trained on three of the five basic jobs in the assembly process, and then teams were composed of people who had received different patterns of initial training. This provided both commonality of skills within teams, since more than one member of each team was competent to do each subtask, as well as moderate heterogeneity, since no one person was initially competent to perform all

the subtasks. This balance both gave teams some flexibility in how they deployed their human resources, and promoted interdependence and cross-training among members.⁴⁴

The natural social forces that sand off the rough edges of teams and polish their compositional center are strong and mutually reinforcing. A good place to start in countering or redirecting those forces is to make sure that interesting differences among members are built into a work team when it initially is composed. But good composition by itself is insufficient. As we will see in chapter 6, it also takes expert coaching to help a diverse group—especially one whose diversity has to do with the personal and group identities of members, not just their task skills—find ways to constructively learn from their differences and exploit them in carrying out the team's work.

Interpersonal Skills

Some people just are not cut out to be team players. They may have a great deal to contribute to the work of an organization, but those contributions are better made as solo performers than as team members. All the teams on which such a person serves have difficulties that appear to stem from his or her disruptive or inappropriate behavior. The person may behave so abrasively as to alienate teammates. Or head off in an independent direction even after the team has decided what it will do. Or escalate conflicts far beyond their actual importance. Or regularly misunderstand what others are asking or suggesting. Or just generally behave as a nuisance, getting in the way rather than pitching in and helping move the work forward.

Without question, there is variation in organization members' interpersonal skills, in their ability to work competently with others (yes, you may call it "emotional intelligence" if you want). This variation is seen in all functions and at all organizational levels. Becoming a senior manager, for example, does not automatically confer on a person interpersonal skills that he or she did not have before. It also is unquestionably true that some basic level of interpersonal skill is required to function well on a team. Therefore, one ideally would compose a team entirely of members who have the requisite levels of interpersonal skill. If some people do not have those skills, they can be offered training to help them acquire such

skills. Sometimes widespread training in interpersonal skills may be required. In organizations that have a long history of designing and managing work for individual performers, most organization members may have mastered well the skills needed for individual work, but more than a few may find that the transition to teamwork stretches their existing skill set.

What can be done with people who do not have the skills needed for competent teamwork and who are unwilling or unable to acquire them? There are only three ways to deal with such individuals when teams are formed. First, keep them at a safe distance so they can do no damage. Indeed, some firms these days try to be rid of them altogether. "Only team players at this company!" is the slogan, as if being a team player were the ultimate measure of one's worth, which it is not. Second, go ahead and put them on teams, install strong leaders to keep things under control, and hope for the best. "Everybody here works on teams. No exceptions!" is the motto, as if everybody were skilled in teamwork, which they are not.

Neither of these alternatives has much to recommend it. The first is wasteful because talent is knowingly withheld from teams. The second is dangerous. Team after team can be sunk by "team destroyers" whose brilliance is exceeded only by their incapacity for collaborative work.⁴⁵ Individuals with marginal task competence are less of a problem. If they persist in misbehaving, the team can afford to do without them. But it is hard to contemplate shunning someone whose task skills are extraordinary.

The only realistic alternative, then, is for a team to find a way to harvest the contributions of those with shaky interpersonal skills, and do so in a way that minimizes the risk to the team and its work. Teammates and the team leader, working together with such a person, often can make significant progress in helping him or her to learn the basics of what teamwork requires. This can be done in the course of the team's regular work and sometimes results in greater progress than would be obtained from his or her participation in even a well-executed interpersonal skills workshop or training course. If peer coaching does not have the hoped-for effect, then the team still has the option of isolating or working around the person so that it can get on with its work relatively unimpeded by his or her disruptions.

Working around a member whose interpersonal skills are viewed as insufficient or inadequate in effect dismisses that person from the team. That is a draconian action, to be taken only after everything else that

members can think of has been tried. Excluding a member from participation in the work of a team can be quite costly to the team itself. The most obvious cost is that the team loses all that the disruptive member has to contribute to the group. That may include some special talents for task work, to be sure, but the loss of the disruptiveness can itself be costly. As noted earlier, group members who violate the norms of a team, who persistently suggest things that everyone else knows is wrongheaded or who behave in ways that are coded as out of line, make their own special contributions to a team. To lose a "deviant" member often is also to lose opportunities for fresh thinking.

A more serious risk of dismissing or ignoring a person whose interpersonal skills seem inadequate is that the attribution itself may be wrong. We all tend to view behaviors that are disruptive or off-target as reflecting some problem with the person who exhibited them. But sometimes those behaviors may actually reflect the perspective of some other group that is not otherwise represented in our own. If, for example, a team consists of five marketers and one engineer who seems constantly to be acting up, that may indeed reflect something that is not quite right about that particular person. It is possible, however, that the person is manifesting an alternative perspective or way of operating that characterizes the other group of which he or she is a member.⁴⁶ The person may be doing that entirely unconsciously, without even being aware that the problems he or she is having in the group actually reflect tensions between that group and the engineering department. If that possibility is not explored, the team runs a real risk of socially or behaviorally dismissing the engineer and, in the process, avoids dealing with intergroup issues that ought to be addressed rather than suppressed. This phenomenon is especially pervasive, and its consequences especially severe, when the other group is central to the identity of the person whose behavior is being questioned—for example, when the person's race, gender, ethnicity, or nationality differs from that of the majority members.

Another potentially destructive misattribution regarding the causes of seemingly unskilled interpersonal behavior derives from a psychodynamic phenomenon known as *splitting*.⁴⁷ In emotionally charged settings, people sometimes deal with their uncertainties and ambivalences about how things are going by unconsciously splitting the positive and negative affect they are experiencing into separate parts, assigning all the former to

one person (the hero) and all the latter to another (the bum). Splitting can be especially pernicious in teams where there are only a few members with distinctive demographic attributes, because those are the ones who are disproportionately likely to be scapegoated, viewed as "unable to work in groups," and assigned most or all of the blame for collective failures. If that bad actor could just be removed, the thinking goes, the team's problems would disappear. The impulse to scapegoat someone when the going gets rough can be quite strong; even worse, the scapegoated member often starts to behave in accord with his or her teammates' expectations—thereby setting in motion a self-fueling spiral whose outcome does no good for either the team or the person.

The final reason teams should move slowly on an impulse to blame and dismiss people whom they experience as disruptive has to do with the topic of this chapter—team structure. One more time, a reprise of the now-familiar refrain: If a team's structure is appropriate, members will encounter fewer interpersonal problems than they would if the team's task, norms, or composition were ill considered. Moreover, the snags a well-structured team inevitably will hit are more likely to be resolved without resort to the kinds of dehumanizing misattributions just discussed.

Putting It Together

Good composition of a work team requires leaders to attend simultaneously to the attributes of the individuals who will make up the team and to the properties of the team as a whole. At the individual level, the main focus should be on ensuring that each member has strong task skills and, as discussed earlier, at least adequate interpersonal skills. At the group level, there are two key considerations. First, the size of the group: It should be small, perhaps even a bit smaller than seems needed to accomplish the work. Second, the mix of members: Team composition should balance between homogeneity and heterogeneity, with special attention to countering the natural social forces that tilt teams toward similarity among members and uniformity of belief, attitude, and behavior.

Herman Melville in *Typee* describes the makeup of the whaleship *Dolly's* crew as follows: "Unfortunately, with a very few exceptions our crew was composed of a parcel of dastardly and mean-spirited wretches, divided among themselves, and only united in enduring without resistance

the unmitigated tyranny of the captain."⁴⁸ With thoughtful attention to how work teams are composed, we should at least be able to do better than that.

BREATHING LIFE INTO A TEAM STRUCTURE

Those who create work teams can do a great deal before a team ever meets to get things set up right, to stack the cards so that competent teamwork will be easy rather than a struggle. The team task can be designed so it engages and motivates members for work on the collective task. Basic norms of conduct that promote active strategy planning within broad but clearly defined limits can be articulated. People with the right skills and experience can be selected for membership on the team, and care can be taken to make sure that both the size of the team and the mix of members are as good as they can be. These structural features comprise what Robert Ginnett and I call the *shell* of the team. A team's shell, like that of an egg, is the shaping structure within which an organism (in the present case, a social system) comes to life. Until an egg is fertilized, it is lifeless; until a team meets for the first time, it is lifeless as well.

Ginnett and I stumbled onto the idea of the shell, and came to appreciate the full importance of that first team meeting, in the course of our research on airline and military flight deck crews.⁴⁹ We went into the research assuming that the behavior of captains made a great deal of difference in how their crews functioned, so we were surprised to discover how many of the determinants of flight crew behavior were already in place when crews first convened and were effectively beyond any captain's control. Both the design of the flying task and crew members' roles are preengineered, for example. Basic norms of conduct on the flight deck are largely dictated by regulatory authorities and airlines' flight standards staffs. The composition of crews is determined jointly by the predefined structure of crew roles and, in most commercial airlines, by a seniority-based bidding system that, in turn, is enforced by a labor contract.⁵⁰

The shell that has evolved for flight deck crews, happily for frequent fliers, is quite sound. What the captain does when the crew first forms is indeed consequential for future team dynamics (see chapter 6). But even if the captain does little more than help members occupy comfortably the

standard structure the crew has been provided, things are likely to go fine. It's like slipping into a suit of one's own size at the clothing shop; all that is required is a little tailoring to perfect the fit. If the captain accepts the basic structure, helps members of his or her team do the same, and refrains from behaving disruptively, the crew is almost certain to get off to a good start.

Some team shells in other kinds of organizations, however, are badly flawed, with the design of the work, expectations about appropriate conduct, and the composition of the team so far off the mark that teams are doomed from the start. In such cases, there is little that any team leader can do to remedy the situation after the group has started its work—it can be difficult, or even impossible, to change the task, to redefine basic norms of conduct, or to adjust the number or mix of members on line and in real time.

That is why it is so important to get the basic structure right, to set things up so the team has every favorable chance of success. But it is only when the team first meets that life is infused into that structural shell. Those who create teams, therefore, have two quite different but equally important responsibilities: to make sure that a team has the best structure that can be provided, and to help members move into that structure and competently launch themselves onto a course of their own. To focus on either one of those responsibilities to the exclusion of the other is to significantly handicap a team even as it begins its life.

VIRTUAL TEAMS: THE END OF STRUCTURE?

It is a fantasy—a tempting and pervasive one, but a fantasy nonetheless—that it is possible to have great teams without the bother of creating enabling team structures. We hope that markets will make hierarchies unnecessary. That we can have networks rather than organizations. That boundaryless social systems can accomplish work efficiently and effectively. And, when some kind of structure actually is needed, that self-organizing processes of the kind celebrated by complexity theory will create them automatically. Both positive models, such as W.L. Gore & Associates, where there are no managers, and horrifying models, such as terrorist groups where there is no visible organization, are held up as previews of the coming of the new organizational age.⁵¹

Now, on wings of ever more powerful, sophisticated, and compact

computational and communications technologies, comes yet another possibility that may allow us to move beyond the necessity to explicitly create team structures: virtual teams. Such teams are becoming increasingly popular these days in forward-looking organizations, especially for knowledge work but also for front-line activities that keep individual team members on the road much of the time, such as sales and service work.⁵² In virtual teams, members interact mainly (and sometimes exclusively) electronically. Virtual teams can be larger, more diverse, and collectively more knowledgeable than those that rely on face-to-face interaction because members do not have to be geographically co-located. And when virtual teams work well, they can bring widely dispersed information and expertise to bear on the team's work quickly and efficiently.⁵³

Do these developments mean the end of structure for work teams? Will managers no longer have to worry about the difficult decisions regarding team design that we have covered in this chapter? I may be proven wrong by the time this book is published (the virtual world is developing faster than I am writing), but I predict the opposite. I am guessing that we will see in the years immediately to come a new millennium version of the "tyranny of structurelessness" that Jo Freeman wrote about in the 1970s and with which I opened this chapter. Virtual teams will be tossed together even more casually than some organizational task forces are these days; it will be unclear who is actually a team member; norms of conduct will be set, if at all, only at the most general level (e.g., "we don't flame each other"); and although the teams will be quite diverse in membership, they will be far too large for members to work together efficiently.

Research evidence about the conditions required for virtual team effectiveness is just now beginning to accumulate.⁵⁴ My reading of that evidence is that the structural conditions that foster the effectiveness of face-to-face teams are just as critical for virtual teams—but with one caveat: *It is much harder to create those conditions in virtual teams.* It is a significant managerial task to create an appropriate team structure—let alone to bring a team to life as a real social system—if members are scattered around the region, the country, or the world. Many practitioners are realizing this, and increasingly are taking the costs in time and travel to bring all members of a virtual team together at least for a face-to-face launch meeting at which the team establishes itself as a group, comes to terms with its task, and sets the basic norms of conduct that will guide

members' interactions online. In my view, that is all to the good, well worth the cost and trouble. But I probably would go even further and ask that members also come together at critical milestones in their team life cycle to debrief, to reflect on ways they might have better used their electronic resources, and to make whatever changes may seem wise in how they will carry out their collective work and operate as a virtual team in the future.

Team structure, as challenging as it can be to get it right and implement it well, will always be with us. That is as it should be, since it is a team's basic structure that provides the platform on which members do all of their work. It always is a good investment to devote time, thought, and energy to making that platform as high and sturdy as it can be—regardless of whether members will be working together around the same table or will be dispersed around the globe, communicating and coordinating their activities electronically.